

TESTIMONY OF

MICHAEL M. ROBERTS

PRESIDENT AND CHIEF EXECUTIVE OFFICER

INTERNET CORPORATION FOR ASSIGNED NAMES AND NUMBERS

Before the

SENATE COMMITTEE ON COMMERCE, SCIENCE AND TRANSPORTATION

Subcommittee on Communications

February 14, 2001

I. INTRODUCTION

Mr. Chairman and members of the subcommittee, I appreciate the opportunity to appear here today and offer a status report on the Internet Corporation for Assigned Names and Numbers (ICANN), which I have served as President and Chief Executive Officer since its formation in November of 1998. I am also happy to report that the "only a few months" assignment, as it was described when I came out of retirement to take it on a little over two years ago, will finally come to an end next month, when I will retire once more, and Stuart Lynn will take over this challenging but interesting task.

Before I report to you on ICANN, I would like to take just a minute to set some context. The reason why there is a need for an ICANN-like organization today is directly traceable to the enormous, worldwide success of the Internet. The Internet's success, in turn, is a product of a sustained commitment by the US government over many years to a public-private partnership among federal research agencies, our pre-eminent research universities, and the energy and entrepreneurial ingenuity of American high technology companies.

Beginning with the Defense Department's sponsorship of the original basic research, and moving through the participation of other agencies such as Energy, NASA, and particularly the National Science Foundation, those of us involved in this process have had at every key point in the evolution of Internet technology, infrastructure, and commercial deployment the kind of US government support that was needed. However, we should also recognize the many contributions of our international partners, which have been essential to the worldwide development and deployment of the network. Indeed, if the Internet was not based on a solid foundation of international partnership, many of the opportunities which it offers for trade, economic development, enhancement of national security and the growth of democratic institutions would not be possible.

The important role of Congress should also be acknowledged. It was my privilege in my former role as a technology policy advocate for higher education to work closely with Mr. Boehlert and Mr. Brown and other members of the House Science Committee in the middle 1980's on legislative programs for support of broader use of the Internet in research and education. It is notable that the High Performance Computing and Communications Act, which President Bush signed in 1991, originated in this Committee.

Hearings such as this, and the recent hearing in the House, continue the constructive tradition of the Congress of encouraging the continued development of a stable, secure and open infrastructure for global commerce and communication.

ICANN itself is a unique entity, but it follows a great tradition of finding and using practical means to address problems that stand in the way of progress. Several years ago, the US government was confronted with the fact that its agency assignments for coordination of Internet activities were seriously lagging the rate at which the Internet was growing, especially in areas related to commercial use. To very much shorten an interesting story, the result of scrutiny of the issues involved was a judgment that the most appropriate solution was to entrust the management of a small set of key technical infrastructure management and coordination responsibilities to the private sector. ICANN was the response of the Internet community to the call for the creation of a private sector, non-profit, global consensus development entity to take over these functions.

In an important sense, the strenuous effort that resulted in the creation of ICANN was the last public service by Dr. Jon Postel, who sadly is no longer with us. His almost thirty-year stewardship of the Domain Name System has left us with a remarkable legacy of selfless devotion to the public interest, along with a basic framework for ICANN's functions that is of important and continuing value.

ICANN was recognized by the US government in November, 1998, by means of a Memorandum of Understanding between the Department of Commerce and ICANN. It was and still is the case that ICANN and its stakeholders are required to earn the trust of the citizens and nations of the world and their governments by demonstrating that private sector consensus management of these functions works efficiently and serves the public interest while promoting opportunities for businesses to engage in the research, development and delivery of network services.

Although ICANN was formed in 1998, we have really been operational for only about 14 months. I think it is fair to say that much has already been accomplished -- indeed, more than some imagined could be done, either in that time or by this entity. For example, there has been a dramatic transformation in the domain name registration market, from a monopoly to an extremely competitive market, with predictably positive impacts on consumers -- including cutting the average price for registration in half. We

now have a well-functioning global dispute resolution system for certain of the most common domain name disputes -- a system that one recent commentator stated was "widely viewed as a model of dispute resolution for the 21st Century." And we are on the verge of introducing real competition at the domain name registry level, a goal that has been fiercely debated and energetically pursued for much of the last decade, but for various reasons never able to be accomplished until the creation of ICANN.

But these achievements, real and important as they are, are only part of the story. We have certainly not yet accomplished ICANN's ultimate goal -- to become a truly effective consensus development body for the entire Internet community in the areas for which ICANN is responsible. We have been forced by events and the speed of Internet time to undertake some complex operational tasks, even though we are still working to complete the basic organizational architecture of ICANN. All the necessary parts are not yet in place. We have certainly not solved the very difficult problem of how to create a global process that is, on the one hand, broadly viewed as fair and effective, but on the other hand, does not erect a procedural, political and legal thicket that makes it impossible to achieve the kind of consensus decision-making that ICANN was created to accomplish.

As a result, no one is really satisfied with the current state of affairs, and rightly so. As it turns out -- and this will be no surprise to any member of this Committee -- achieving global consensus is a difficult task, especially on issues as complex and important as those which prompted the creation of ICANN. There are important parts of the Internet community -- country code registry operators, address registry operators, root server operators, and the general user community -- where we have not completed the discussions that will formalize their relationships with or within ICANN. Even those elements of the construction that appear to have been completed, such as ICANN's three Supporting Organizations, need refinements of various kinds; any structure that is, as ICANN was, the product of a series of compromises is not likely to be perfect at first creation.

And so, while we have attempted to be responsive to the important operational objectives that formed much of the impetus for the creation of ICANN, we have also worked very hard -- and we continue to work hard -- to assemble a complete working organization for the development of global consensus on these issues, and to ensure that

all the stakeholders in the Internet community have an appropriate place in, and the ability to have their voices heard in, the ICANN process.

I am frequently asked, "Why is there so much noise around ICANN? How can you get any work done over there?" My response is that ICANN is noisy by design. We are intended to be the forum in which interested parties - some might characterize them as combatants - have the opportunity to advance multiple futures for the domain name and address system, and have those competing and frequently contradictory futures merged into one satisfactory solution. By definition, it will be noisy, and messy, and sometimes slow, and frequently contentious, but if it works -- and the jury is still out, although I am reasonably optimistic -- it may well be a useful model for other global issue resolution mechanisms.

II. BACKGROUND

For much of its formal history, which begins in 1973 with roots stretching into the 1960s, the functions of ICANN were performed by one computer scientist, Jon Postel, under a research contract to the US Defense Advanced Research Projects Agency (DARPA). During the mid 1990s, as the Internet emerged as a potent commercial force in the telecommunications environment, it became clear that such functions needed to be institutionalized. Dr. Postel participated in attempts to achieve that goal beginning as far back as 1995. In the midst of the effort in the late 1990's that led to the creation of ICANN, Dr. Postel unexpectedly passed away. ICANN was formed to privatize, institutionalize and internationalize the functions that Dr. Postel performed so ably for so long.

A. The Formation of ICANN

ICANN is a non-profit private sector organization with a 19-member international volunteer Board of Directors¹ drawn from a set of specialized technical and policy advisory groups, and from an online voting process of Internet users worldwide. Through a series of Supporting Organizations, Advisory Committees and Working Groups,² it functions as a consensus development body for certain technical and administrative management issues related to the name and address functions of the Internet.

ICANN is the end result of an extensive policy development process, both within the United States Government and within the global Internet community. During 1997 and 1998, under the leadership of the US Department of Commerce, a framework for private sector management of the Internet's Domain Name System (DNS) and Address System was developed and put into writing in the form of a policy document known as the White Paper.³

The White Paper, which was issued in June 1998, proposed that the private sector undertake management of these functions through the formation of a private, non-profit corporation, and it outlined the substantive responsibilities of the new organization and a number of guiding principles for its work. Following several months of public meetings and dialogue in the summer of 1998, during which the White Paper framework was turned into a specific charter and set of Bylaws, ICANN was incorporated in September of that year, and was recognized by the U.S. Government in November 1998, in the form of a two-year Memorandum of Understanding/Joint Project Agreement between the Commerce Department and ICANN. The MOU has subsequently been amended twice and currently has a term expiring on September 30, 2001.⁴

¹ Short biographies of the directors can be found at <http://www.icann.org/general/abouticann.htm>

² An organizational chart of ICANN and its constituent units is attached to this testimony.

³ The White Paper can be found at <http://www.icann.org/general/white-paper-05jun98.htm>

⁴ The full text of the MOU/JPA can be found at <http://www.icann.org/general/icann-mou-25nov98.htm>

B. ICANN Responsibilities

The White Paper identified four principal areas of responsibility for the new private sector consensus organization:

- Coordination of the Internet Domain Name System
- Overseeing operation of the authoritative root server system
- Coordination of the Internet Protocol (IP) Address space
- Coordinating the assignment of Internet technical parameters

As recognized in the White Paper, these four functions were broadly seen by the global Internet community as requiring coordinated action to assure the smooth and reliable operation of the Internet.

C. Guiding Principles for ICANN

The White Paper identified four principles that it described as critical to the success of an entity such as ICANN: stability; competition; private, bottom-up coordination; and representation.

1. Stability is perhaps the easiest to understand. The US Government was seeking to extract itself from what it had concluded was no longer a proper role for the US Government -- the funding of private contractors by research agencies to manage important technical aspects of the global Internet name and number address system -- but only in a way that did not threaten the stability of the Internet. As the White Paper said, and as seems obvious, "the stability of the Internet should be the first priority of any DNS management system." If the DNS does not work, then for all practical purposes for most people, the Internet does not work. That is an unacceptable outcome, and thus everything that ICANN does is guided by, and tested against, this primary directive.

2. Competition was also an important goal set forth in the White Paper, which stated that "[w]here possible, market mechanisms that support competition and consumer choice should drive the management of the Internet because they will lower costs, promote innovation, encourage diversity, and enhance user choice and satisfaction." Competition in the registration of domain names is theoretically possible at both the registry (or wholesale) level, and at the registrar (or retail) level. Increasing competition at the retail level involves only allowing multiple providers of registration services to add domain name registrations to registry databases; as a result, that

objective can be accomplished without major stability concerns. For this reason, adding new competition at the retail level was the first substantive goal that ICANN quickly accomplished after its formation. By contrast, the introduction of competition at the new registry (or wholesale) level requires the introduction of additional Top Level Domains into the namespace, and thus does raise potential stability issues of various kinds. As a result, and given its prime directive to protect stability, ICANN has moved forward in this area in a prudent and cautious way, consistent with recommendations from many constituencies with a stake in the Internet.

3. A third White Paper principle was **private sector, bottom-up consensus development**, and the entirety of ICANN's processes are organized around this principle. ICANN is a private-sector body, and its participants draw from the full range of Internet stakeholder organizations, from business entities to non-profit organizations to academic institutions to individual Internet users. Its policies are the result of the complex, sometimes cumbersome interaction of all these actors in an open, transparent, sometimes slow and sometimes contentious progression from individuals and particular entities through the ICANN working groups and Supporting Organizations to ICANN's Board, which under its bylaws has the principal role of recognizing consensus as developed below, rather than imposing it from above. Like democracy, consensus is far from a perfect system, but it is an attempt, and the best way we have yet been able to devise, to achieve globally acceptable policies without the coercive power of governments.

4. Finally, the fourth core principle on which ICANN rests is **representation**. A body such as ICANN can only plausibly claim to operate as a consensus-development organization for the Internet community if it is truly representative of that community. The White Paper called for ICANN to "reflect the functional and geographic diversity of the Internet and its users," and to "ensure international participation in decision making." To satisfy these objectives, all of ICANN's structures are required to be geographically diverse, and the structures have been designed to, in the aggregate, to provide opportunities for input from all manner of Internet stakeholders. This is an extremely complicated task, and we are not yet finished with the construction phase; indeed, we have just initiated a Study Committee chaired by Carl Bildt, the former Prime Minister of Sweden, to oversee a new effort to find a consensus approach to obtaining input from and

providing accountability to the general Internet user community, which might not otherwise be involved in or even knowledgeable about ICANN and its activities. This is a formidable challenge, given that there are an estimated 400 million Internet users around the world in over 200 countries -- a number that has been growing at 80-100% per year since 1988.

We have also undertaken a number of other organizational tasks necessary to ensure that ICANN is fully representative of the entirety of the Internet community. This is hard work, and there is more to do to get it done right.

III. ICANN ACCOMPLISHMENTS TO DATE

The tasks assumed by ICANN in the Memorandum of Understanding were of two general kinds. The first were related to completion of its organizational structure, particularly its three specialized Supporting Organizations (for domain names, technical protocols, and IP Addresses -- the numeric identifiers used in Internet routing), and its fourth component, known as "At Large" (which is intended to provide a vehicle for input and participation by the full range of Internet users in ICANN's work). The second set of tasks were related to specific problems that had arisen as a result of the rapid growth and commercialization of the Internet in the middle 1990s.

Obviously, ICANN is still a work in progress. Nevertheless, it has already made remarkable progress in the short span of little more than two years.⁵ In the following portions of the testimony, I describe our work on four specific tasks - the enhancement of the Internet's Root Server System; introduction of retail competition in domain name registrations for .com, .net, and .org; adoption of a non-judicial mechanism for resolving certain disputes over the registration of domain names; and introduction of new Top Level Domain Name Registries to provide "wholesale" competition. Because staff has indicated that the Committee has a special interest in the Root Server System, I will begin with that subject.

A. Enhancement of the Security and Reliability of the Root Server System

A.1 Functioning of the Domain Name System. In recent years, the domain-name system (DNS) has become a vital part of the Internet. The function of the domain name system is to provide a means for converting easy to remember mnemonic domain names into the numeric addresses that are required for sending and receiving information on the Internet. The DNS provides a translation service that permits Internet users to locate Internet sites by convenient names (e.g., <http://www.senate.gov>) rather than being required to use the unique numbers (e.g., 156.33.195.33) that are assigned to each computer on the Internet.

⁵ For the first two annual summaries of progress provided to the Department of Commerce, see First Status Report at <http://www.icann.org/general/statusreport-15june99.htm>; Second Status Report at <http://www.icann.org/general/statusreport-30jun00.htm>.

The Internet engineering community devised the DNS in the early 1980s.⁶ One of the Internet's prominent engineers, Dr. Jon Postel (the creator of the IANA function that preceded ICANN, and the principal force behind the creation of ICANN) took on responsibility for coordinating a decentralized system of computers throughout the Internet to implement the DNS. These computers are organized in a hierarchical manner, with "root nameservers" at the highest level that point to nameservers for top-level domains (e.g., .gov), that in turn point to nameservers for second-level domains (e.g., senate.gov), and so on. In all there are 253 top level domains, of which the greatest number are assigned to the national, or "country code," top level domains.

Upon the deployment of this new system in 1985, Internet users worldwide could point their computers to the root nameservers, and use them to receive the translation services (i.e. from names to numbers) that the DNS provides. The system is highly redundant and decentralized, consisting of almost 100,000 nameservers arranged in a topologically and geographically distributed system. It has repeatedly demonstrated its technical resilience and robustness, including during last year's Y2K event during which the system functioned smoothly without interruption.

As a first step in deploying the DNS nameserver system, Dr. Postel arranged for voluntary operation of the root nameservers by a group of expert and trusted individuals and organizations throughout the world, who each volunteered to operate a root nameserver. This group now numbers nine organizations, plus the U.S. Government; they operate the thirteen root nameservers on a completely voluntary, free-of-charge, and public interest basis. The following map and chart show the identities and locations of the organizations operating the DNS root servers:

⁶ The DNS replaced an earlier, smaller capacity translation mechanism known as the "hosts.txt" system.

At lower levels in the DNS hierarchy (for example .com), the operators of the nameservers and the associated registries have received compensation, first by governmental subsidies in the late 1980s and early 1990s and then, beginning in the mid-1990s, by charging those who wished to register domain names within the system. The root nameserver system itself, however, has always been operated on a voluntary basis and without user fee (or even government subsidy, though the U.S. Government does contribute by operating some of the thirteen root nameservers). As a result, the system has become broadly accepted by Internet users worldwide as an integral feature of the Internet.

A.2 U.S. Government Policy Concerning the Root Server System. As the Internet has evolved from a system for research conducted under U.S. Government sponsorship to an essential medium for global commerce, the need for a secure, stable, and reliable DNS root nameserver system coordinated according to the needs of the Internet community has also grown. The White Paper reflected a broad consensus within the Internet community when it said, "coordination of the root server network is necessary if the whole system is to work smoothly. While day-to-day operational tasks, such as the actual operation and maintenance of the Internet root servers, can be dispersed, overall policy guidance and control of the TLDs and the Internet root server system should be vested in a single organization that is representative of Internet users around the globe."

In the ICANN MOU, the US Government represented that it would "undertake, in cooperation with IANA, NSI, the IAB, and other relevant organizations from the public and private sector, a review of the root server system to recommend means to increase the security and assure professional management of the system. The recommendations of the study should be implemented as part of the transition process; and the new corporation should develop a comprehensive security strategy for DNS management and operations."

Subsequently, one of the first Advisory Committees established by ICANN was the Root Server System Advisory Committee (RSSAC), chaired by ICANN Board member Jun Murai of Keio University in Japan. Professor Murai is also responsible for the operation of the root nameserver located in Tokyo. All root nameserver operators are members of the RSSAC, which also includes technical experts from the Internet Engineering Task Force (IETF).

The RSSAC has been working diligently since ICANN's creation to evaluate and improve where necessary the security and reliability of the root nameservers. In its last report, at ICANN's public meeting in Yokohama in July of last year, it described the results of its efforts, which basically involve the evolution in the near future of the current root nameserver system structure to one in which a "dedicated primary" server, rather than one of the 13 operational root servers, is responsible for distributing updated root zone files to the publicly accessible root nameservers in a secure, reliable and robust system transparent to users. When implemented, this will be a major improvement in the security and reliability of the root nameserver system, and therefore of the DNS and the Internet.

A.3. Formalization of Arrangements for Operation of the Root Nameservers

In addition, ICANN has been working on formalizing the legal relationships under which the various organizations have operated the individual DNS root nameservers. As described above, since the initial deployment of the DNS the root nameservers have been operated under the voluntary arrangements originally made by Dr. Jon Postel. After ICANN was established, some additional formality was introduced by the participation of the operators in the RSSAC, and in mid-1999, ICANN and the National Institute of Standards and Technology entered into a Cooperative Research and Development Agreement under which the U.S. Government is participating in the RSSAC's work toward enhancing the stability and security of the root nameserver system. As part of this effort, ICANN is near the completion of agreements with the organizations operating the individual root nameservers, with the goal of mutually recognizing in an appropriate way each other's obligations and responsibilities to protect the stability of the DNS and the Internet. We are well along in those discussions and I expect they will be completed in the near future.

A.4 Administration of Changes to the Root Server System.

There has been, and continues to be, some confusion about the current and proposed procedures for coordination and administration of changes to the files contained in the root server computers.

Currently, the root nameserver operators follow the convention that one of the operators, Network Solutions, Inc. (NSI), is responsible for implementing edits to the “root zone” file that designates the top-level domains in the DNS. Under agreements among ICANN, the U.S. Government, and NSI, ICANN (through IANA, now absorbed into ICANN), sends documentation for needed changes to the root zone file to the U.S. Department of Commerce, which directs Network Solutions to implement them by editing the authoritative root zone file. By convention among the RSSAC’s root nameserver operators, that file is loaded twice daily into all thirteen DNS root nameservers.

ICANN, through the RSSAC and through its soon-to-be-completed agreements with the root server operators, is already playing an important role in facilitating a more structured understanding among these most critical participants in the DNS. As a result, the very informal arrangements that have served us well in the past are in the process of a transition to a more transparent but still collegial and consensus-based structure that we believe will continue this outstanding record of service into the future.

B. The Introduction of Retail Competition

A very important impetus for the formation of ICANN was the perception that the name registration market was not competitive, and as noted above, the introduction of competition was an important goal outlined in the White Paper. Thus, as one of its very first actions, ICANN created an accreditation system for competitive registrars and, pursuant to its agreements with NSI, gave those new competitors access to the NSI-operated registries (specifically, .com, .net and .org).

When ICANN was formed, there was only a single registrar (NSI) for .com, .net, .org, and everyone had to pay the single price for the single domain name product that sole registrar offered: \$70 for a two-year registration. There are now over 180 accredited registrars, with more than half of those actively operating, and you can now register a domain name in the .com, .net, and .org registries for a wide range of prices and terms - some will charge zero for the name if you buy other services, while others will sell you a ten-

year registration for significantly less than the \$350 it would have cost pre-ICANN (even if it had been available, which it was not). While there are no precise statistics, in part because the market is so diverse, a good estimate of the average retail price today of a one-year domain name registration in the NSI registries is probably \$10-15 -- or less than half the retail price just 18 months ago.

As another illustration of the dramatic changes over the last year, NSI's share of the registration market for the .com, .net and .org TLDs has fallen from 100% at the time of ICANN's creation to less than 40% of new registrations in those TLDs today -- a market share drop of more than half in just over a year.

There are still issues that must be dealt with in this area; some registrars appear not to have lived up to their contractual commitments, and ICANN needs to ensure that they do. And indeed, there may be more registrars than the market will support in the long term; 94% of all registrations come from the 10 largest registrars, with the other 80 or 90 active registrars sharing the other 6%. Name registration is quickly becoming a commodity business, and a commodity business, with commodity margins, will probably not support 100 vigorous competitors. We are already starting to see some companies wishing to leave the business, and we need to make as sure as we can that those departures do not impair the ability of consumers and businesses to rely on names they have registered, and that departures or even failures do not generate unreliability or other forms of instability in the namespace itself. While these issues must be dealt with, I think it is widely recognized that ICANN has been very successful in changing the retail name registration market from a monopoly market to a highly competitive market.

C. Creation of a Cost-Effective, Efficient Dispute Resolution System.

Another significant accomplishment has been the creation of the Uniform Dispute Resolution Policy (UDRP), a way to quickly and cheaply arbitrate certain domain name disputes. While domain names themselves cannot be trademarked, it is certainly possible for domain names to be confusingly similar to a trademarked name, or in other ways to be inappropriately used by someone for illegitimate means. Since trademark and other intellectual property rules differ from country to country, enforcing those rights is complex and expensive.

One of the policies that was generated from the ICANN bottom-up process early on was the need for a simple procedure to resolve the clearest and most egregious violations on a global basis. The result, after considerable work in a variety of ICANN forums, is the UDRP, which one commentator recently noted is "widely viewed as a model of dispute resolution for the 21st Century." The UDRP is limited to certain very specific claims, is intended to require only about \$1,500 in costs and 45 days to invoke, and is required to be included in all name registration contracts by all ICANN-accredited registrars, thus providing the basis for global uniformity in the resolution of this particular class of domain name disputes. Even though the UDRP is non-binding (either party may take the dispute to court after an unfavorable UDRP decision), it appears that has happened in only a few dozen cases out of over 2,000 decisions to date.

The UDRP is, I would submit, another very positive accomplishment of ICANN during its short existence to date. As of this writing, parties interested in further refinement of the UDRP are already studying its design for possible revisions.

D. The Introduction of New TLDs

D.1 Background. This brings us to the current effort to introduce competition at the registry (or wholesale) level of the domain name market. ICANN was able to create retail competition relatively quickly after its creation, and this has produced the expected benefits -- lower prices, more consumer choice, and innovation. But the introduction of wholesale competition, because it involves actually expanding the structure of the namespace, presented and continues to present more risks. While most Internet engineers believe that some number of additional TLDs can be added without serious risks of instability, there is considerable uncertainty about how many could be added without adverse side effects, and very few engineers have been willing to absolutely guarantee that there was zero risk of instability. Given the increasingly critical role the Internet now plays in everyday commercial and personal life, the almost uniform consensus in the community was to be cautious and prudent in this process.

For example, the White Paper asserted that "expansion of gTLDs [should] proceed at a deliberate and controlled pace to allow for evaluation of the impact of the new gTLDs and well-reasoned evaluation of the domain space." In addition to concerns about the technical stability of the Internet, many were concerned about potential costs that rapid

expansion of the TLD space might impose on business and consumers. The World Intellectual Property Organization, which conducted a study of intellectual property issues in connection with the DNS at the request of the United States Government, concluded that new gTLDs could be introduced if done "in a slow and controlled manner that takes into account the efficacy of the proposed measures in reducing existing problems." The Protocol Supporting Organization of ICANN (made up of the Internet Engineering Task Force and other Internet engineering and communications protocol development bodies) said it saw no technical problems with the introduction of a "relatively small" number of new TLDs.

In fact, every entity or organization without an economic stake in the answer that has examined this question has recommended the same thing: a "small" or "limited" or "prudent" number of new TLDs should be tried first, as a sort of proof of concept or experiment. Once this "limited" number of new TLDs was introduced -- and the suggested numbers roughly ranged from 1 to 10 -- and assuming there were no adverse side effects, then additional TLDs could be introduced if there was consumer demand for them.

D.2 ICANN Process. Because ICANN is a consensus development body that relies on bottom-up policy development, the issues of whether and how to introduce new gTLDs were first taken up by the Domain Name Supporting Organization (DNSO), the ICANN constituent body responsible for name policy issues. The DNSO organized a Working Group, which recommended that a small number (6-10) of TLDs be initially introduced, and that the effects of that introduction be evaluated before proceeding further. That recommendation was forwarded to the Names Council, the executive body of the DNSO, which reviewed the Working Group recommendation and public comments on it, and recommended to the ICANN Board that it establish a "policy for the introduction of new gTLDs in a measured and responsible way." The Names Council suggested that "a limited number of new top-level domains be introduced initially and that the future introduction of additional top-level domains be done only after careful evaluation of the initial introduction."

Consistent with the ICANN bylaws, the ICANN Board accepts the recommendations of Supporting Organizations if the recommendations meet certain minimal standards designed to ensure that they truly represent a consensus position. Thus, the Names

Council recommendation was published for public comments, and following the receipt of numerous public comments, the ICANN staff in June 2000 issued a Discussion Draft seeking public comments on a series of questions intended to lead to the adoption of principles and procedures to be followed in a "measured and responsible introduction" of a limited number of new TLDs.⁷ Following several thousand additional public comments, and considerable discussion at a public meeting in Yokohama in July 2000, the ICANN Board adopted a series of resolutions instructing its staff to begin the process of accepting applications for a "proof of concept" for the introduction of new TLDs.⁸

D.3 Criteria for Evaluating Applications. In early August, ICANN posted a detailed discussion of the new TLD process it proposed to follow⁹, and in mid-August a detailed set of Criteria for Assessing TLD Proposals.¹⁰ These nine criteria have been constant throughout this process, and so they bear repeating here:

a. The need to maintain the Internet's stability.

This speaks for itself. ICANN's overriding obligation is to protect the stability of the Internet, and all other objectives are subordinate to that. Thus, any proposal that could be shown to threaten this stability (other than any risk inherent in any new TLD introduction) was obviously unacceptable.

b. The extent to which selection of the proposal would lead to an effective "proof of concept" concerning the introduction of top-level domains in the future.

This too is largely self-explanatory. The effort here was not to find the "best" application, however that might be measured, but to ask the community to offer up a set of options from which ICANN could select a limited number that, taken in the aggregate, would satisfy the evaluation objectives of this proof of concept. This is exactly the same approach that ICANN had previously taken in the introduction of competitive registrars, and

⁷ See generally *ICANN Yokohama Meeting Topic: Introduction of New Top-Level Domains*, at <http://www.icann.org/yokohama/new-tld-topic.htm>.

⁸ See *Resolutions of the ICANN Board on New TLDs*, at <http://www.icann.org/tlds/new-tld-resolutions-16jul00.htm>

⁹ See *New TLD Application Process Overview*, at <http://www.icann.org/tlds/application-process-03aug00.htm>

¹⁰ See *Criteria for Assessing TLD Proposals*, at <http://www.icann.org/tlds/tld-criteria-15aug00.htm>

which had worked so well there. The addition of multiple registrars to the NSI registries required the creation of new interface software, since before this time only one registrar had been able to direct new entries in those registries. Thus, there was some experimental effort required to make sure that the software was ready for use by a larger number of simultaneous registrars. ICANN first created a "test-bed," asked for expressions of interest from the community, and accredited only five new registrars for a period of a few months, while they and NSI worked out the bugs in the interface software. As soon as the test-bed was completed, ICANN accredited larger numbers of registrars, now exceeding 180.

Here, the concept is similar: from options offered up from the community, create a limited number of new TLDs to ensure that the DNS can accept, both technically and practically, these additions without impairing stability in any way. Once that is proven, additional TLDs can be created as appropriate.

c. The enhancement of competition for registration services.

Obviously, this is the principal reason for adding new TLDs, so one criterion for determining which applications to accept initially is how effective they are likely to be in creating new competition for the NSI registries. Of course, competition takes many forms; here, one form would be analogous to .com -- a global, unrestricted registry focusing on business. To compete in this way requires not only desire, but the capacity to effectively compete with a market participant that already has high brand awareness, a very significant marketing budget, and a large installed base of registered names which will produce some level of renewals more or less automatically. To compete successfully on a global basis under these circumstances requires a significant capital investment, very significant technical expertise (running a database of several million names that gets hundreds of queries every second is a complicated matter), and a substantial marketing budget to build the kind of brand equity that will be necessary to compete effectively with, for example, .com.

Another way to introduce competition into the wholesale part of the market is to offer a different kind of product -- not a global unrestricted domain, but various kinds of limited or restricted registries that might appeal to specific different sectors of the market. To use a television analogy: narrowcasting instead of broadcasting. Here, capital and

marketing expenses may be lower, but other kinds of service characteristics may be more important.

ICANN's purpose with this criterion was to invite a broad range of competitive options, from which it could select a menu that, taken as a whole, would offer a number of different competitive alternatives to consumers of domain name services.

d. The enhancement of the utility of the DNS.

In addition to competition, one must reasonably consider the practical effects of the introduction of new TLDs. The names registered in the DNS are intended to be used by people, and sound engineering requires that human factors be taken into account, so that confusion, recognition difficulties, and the like do not impair the DNS's ease of use.

e. The extent to which the proposal would meet previously unmet types of needs.

If it is assumed that the DNS should meet a diversity of needs, it would be a positive value if a proposed TLD appeared to meet any previously unmet needs of the Internet community.

f. The extent to which the proposal would enhance the diversity of the DNS and of registration services generally.

Here, what was sought was diversity of all kinds, in the hopes of creating the broadest possible -- and thus most instructive -- experiment within the limitations recommended (i.e., a small number of new top level domains). So, the published criteria encouraged the submission of proposals for different kinds of TLDs (open or closed, non-commercial or commercial, personal or business-oriented, etc.) The criteria also sought diverse business models and proposals from different geographic regions, for the same reasons.

g. The evaluation of delegation of policy-formulation functions for special-purpose TLDs to appropriate organizations.

For those proposals that envisioned restricted or special-purpose TLDs, this criterion recognized that development of policies for the TLD would best be done by a "sponsoring organization" that could demonstrate that it would include representative participation of all segments of the communities that would be most affected by the TLD.

Thus, with this class of application, the representativeness of the sponsoring organization was a very important criterion in the evaluation process.

h. Appropriate protections of rights of others in connection with the operation of the TLD.

Any new TLD is likely to have an initial "land rush" when it first starts operations as people seek the most desirable names. In addition, every new TLD offers the potential opportunity for cybersquatting and other inappropriate name registration practices. This criterion sought information about how the applicant proposed to deal with these issues, and also how it proposed to provide appropriate mechanisms to resolve domain name disputes.

i. The completeness of the proposals submitted and the extent to which they demonstrate realistic business, financial, technical, and operational plans and sound analysis of market needs.

Finally, this criterion simply emphasized that, since the effort was a "proof of concept," the soundness and completeness of the application and the business plan would be important elements of the selection process. This was not intended to be an experiment in how well the DNS or the Internet could survive the business failure of a new TLD operator, nor how businesses and consumers might suffer from a failure. It was also not intended to be clairvoyant with regard to the outcome of any particular proposal. Thus, to the extent possible and consistent with other goals, the Board favored those applications that appeared to have the soundest business plans, and were based on the most realistic estimates of likely outcomes.

D.4 The Application Process and Fee. The application process required the filing of a detailed proposal speaking to all the criteria outlined above. It recommended that applicants retain professional assistance from technical, financial and management advisers, and lawyers. And perhaps most controversially, it required a non-refundable application fee of \$50,000. A brief explanation of this particular requirement may be useful.

ICANN is a self-funding organization. It has no capital, and no shareholders from which to raise capital. It must recover its costs from the various constituent units that

benefit from ICANN's processes and procedures; today, those costs are borne by address registries, name registries, and registrars. Its annual expenditures to date have been in the \$4-5 million range, covering employee salaries and expenses (there are now 14 employees), and a wide range of other expenditures associated with operating in a global setting in an open, transparent, bottom-up consensus based manner.

Thus, there was no ready source of funds to pay for the process of introducing new TLDs, and the ICANN Board determined that this, like all other ICANN activities, should be a self-funded effort, with the costs of the process borne by those seeking the new TLDs. At that point, ICANN estimated the potential costs of this process, including the retention of technical and financial advisers, legal advice, the logistics of the process, and the potential cost of litigation pursued by those not satisfied with the process or the results. While obviously all these elements were highly uncertain, based on its best judgment of how many applications were likely to come in and what the likely costs would be, ICANN established a \$50,000 fee.

As it turns out, there were more applications than expected, and thus the absolute costs of processing and reviewing them were higher than expected; about half the application revenues have already been used to cover costs of the process to date, with considerable work left to do and still with the potential for litigation at the end of the process. To date, it appears that the fact of more applications and higher costs of review and evaluation than expected have cancelled each other out, and so it appears that the fees adopted were about right in creating the funds necessary to carry out this process.

I know there have been complaints by some that they were foreclosed from this process because they simply could not afford the \$50,000 application fee, and I am sympathetic to these concerns. But there are three practical responses that, in my view, make it clear that this is not a fair criticism of the process. First, the process had to be self-funding; there simply was no other option, since ICANN has no general source of funds. Based on costs to date and those projected, it certainly does not seem that the fee was set too high. While there are still application fee receipts that remain unspent, the process is not over, and it has already consumed half of the fees collected. Second, and as importantly, it is highly unlikely that any individual or entity that could not afford the application fee would have the resources to be able to operate a successful and scalable

TLD registry. The capital and operating costs of even a small registry are considerable, and especially if the goal is to operate a registry that charges low or no fees for name registrations (many of the persons and entities advancing this particular complaint are non-profit or public interest bodies), those fees would not likely cover the costs of operation, much less the necessary start-up and capital costs. Of course, it is possible that, if an organization that would otherwise have difficulty managing the costs of operating a TLD registry were in fact awarded a new TLD, it might be able to raise the funds through subsequent contributions or grants or the like, but this leads us directly to the third point.

This effort was not a contest to find the most qualified, or the most worthy, or the most attractive for any reason of the various applicants. ICANN should not be in the business of making value judgments. What ICANN is about is protecting the stability of the Internet and, to the extent consistent with that goal, increasing competition and competitive options for consumers of domain name services. Thus, what ICANN was doing here was an experiment, a proof of concept, an attempt to find a limited number of appropriate applicants to test what happens when new TLDs of various kinds are added to the namespace today -- a namespace that is vastly different in size and in application than that which existed more than 15 years ago when the first seven global TLDs were created.

Because this was a proof of concept, the emphasis was on diverse business models, technical capacity, and diversity of geography and focus -- and not on some weighing of the relative merits, however measured, of the applicants. Indeed, a serious attempt was made to avoid otherwise normal business risks, such as limits on capital or other resources, so that foreseeable likely business failures did not interfere with the data collection and evaluation process of this experiment. Thus, it would have been impossible to accept any application which relied on the mere hope of obtaining funding if an application was accepted, and indeed, several of the applicants were not selected in the evaluation process at least in part on just on that point.

Under these circumstances, it was not appropriate to encourage applications by those with limited resources, since those limitations would almost certainly result in their not being selected. Thus, setting the fee to recover expected costs, without regard to the effect it had on applications, seemed then (and seems today) the logical approach. Once this experiment is over, and assuming it demonstrates that adding new TLDs in a

measured way does not threaten the stability of the DNS or the Internet, I would hope that processes could be developed to both expedite and significantly reduce the cost of new TLD applications or, at a minimum, to deal with special cases of TLDs with very limited scope, scale and cost.

D.5 The Evaluation Procedure. Forty-seven applications were submitted by the deadline established; three of those were withdrawn for various reasons, and the remaining 44 were then published on ICANN's website, open to public comments, and subjected to an extensive evaluation, applying the criteria set forth in the various materials previously published by ICANN. More than 4,000 public comments were received. The applications and the public comments were carefully reviewed by technical, financial and legal experts, and the result of that evaluation -- a 326-page staff report summarizing the public comments and the staff evaluation -- was itself posted on the ICANN website for public comment and review by the Board of Directors of ICANN.¹¹ Another 1,000 public comments were received on the staff report. The Board, of course, had access to the applications and the public comments as they were filed, and was kept generally informed as to the process of the evaluation.

There has been some criticism of the fact that the full staff evaluation was not available to the public -- and thus to the applicants -- until November 10, only days before the actual Board meeting. Obviously, it would have been much better to produce this earlier, and we tried to do so. But in fact the timing of the release of the staff report was largely the product of the bottom-up process that ICANN follows to generate consensus. An important ingredient in the staff evaluations was the substance of the voluminous public comments produced in the month after the applications were posted. ICANN's job is to identify consensus, and thus input from the community is a critical part of any Board decision. Getting that community input, considering it, and completing the technical and financial evaluations was a massive job.

In one sense, it would have been preferable to have issued the staff report earlier. But on the other hand, doing that would have required shortening the period that the public had to make comments that would be summarized in the report. In fact, in the six days

¹¹ See *Report on New TLD Applications*, at <http://www.icann.org/tlds/report>

between the posting of the report and the Board meeting, ICANN received more than 1,000 additional public comments on the staff report, many from the applicants responding to the evaluation of their particular application. The ultimate question is whether the Board got sufficient timely information on which to base its selection decisions, bearing in mind the objective of the exercise. I believe it did.

At its Annual Meeting in Los Angeles in November 2000, the ICANN Board devoted nearly all of the standard public forum day immediately preceding the Board meeting to the new TLD issue, with presentations by the staff of their findings, public comments, and short presentations from the applicants. Another point of criticism by some has been the short time -- three minutes -- allowed during this public forum for presentations by each of the applicants, but oral presentations were never intended to be the sole or primary source of information for the Board. Voluminous applications (with many hundreds of pages) had been filed by each applicant; many of them had received and answered clarifying questions from the staff; and many of them had provided additional material by filing material on the ICANN public comment page (every one of the 5,000+ comments was read by ICANN staff). The Board had access to the applications and to the staff evaluations well ahead of the public Board meeting at which the applications were reviewed. The opportunity to make a presentation at the public forum was simply the final step in an extensive process, available so that any last-minute questions could be asked or points made.

Since there were 44 applicants, nearly all of whom wished to speak, and since the time available for the applicants (given the other parts of the community who also wished to be heard) was limited to about two hours, three minutes was simply all the time available. Most used it wisely, pointing out the particular strengths of their applications.

Some disappointed applicants have also complained that ICANN staff refused to talk with them, or let them respond to concerns raised by their applications. This is not accurate; what ICANN staff refused to do is have *private* conversations with the applicants, and this derives from the very nature of ICANN as an entity. ICANN is a consensus development body, not a regulatory agency; its decisions are intended to reflect consensus in the Internet community, not simply the policy preferences of those who happen to sit on its Board at any given moment. For this process to work, the vast bulk of ICANN's work

must be transparent to the public, and so with very rare exceptions (such as matters dealing with personnel issues), everything ICANN does it does in public. (In fact, one applicant withdrew its application because of its unwillingness to allow significant material in the application to be posted on ICANN's website.) If the public was going to have a real opportunity to comment on the applications, the applications themselves needed to be public, and any substantive discussion of them had to be public as well.

In an effort to help this process, and still get questions answered, ICANN staff frequently took email or other private questions, reformulated them to make them more generically useful, and then posted them on the website as FAQs. In addition, staff encouraged applicants to post any information they wished on the public comment pages, where it would be read by ICANN staff, the ICANN Board and also by any interested observer. What staff would not do, and this was evidently very frustrating to many of the applicants that had not previously had any experience with the open structure and operations of ICANN, was to have private substantive discussions with the applicants.

It is easy to understand this frustration, especially for those disappointed applicants who had not previously participated in the ICANN process and, as a result, did not understand what ICANN is and how it operates and thus were surprised at the transparency of the entire process. Still, it is hard to see how any other process could have been followed consistent with ICANN's consensus development process. Without public access to the entirety of the information about each applicant and each application that was available to the Board, the Board would not have had the benefit of public comments on some (often significant) factors, and it would have been hard to justify its selections as deriving from a consensus development process.

D.6 The Selection Process. To understand the selection process, we must go back to first principles. The goal here was not to have a contest and pick winners; it was not to decide who "deserved" to have a new TLD; it was not even to attempt to predict the kind or type of TLDs that might get public acceptance. The goal, articulated plainly from the beginning of the process more than a year ago, was to identify from suggestions by the community a *limited* number of diverse TLDs that could be introduced into the namespace in a prudent and controlled manner so that the world could test whether

the addition of new global TLDs was feasible without destabilizing the DNS or producing other bad consequences.

This was not a race, with the swiftest automatically the winner. It was a process that was intended to enable an experiment, a proof of concept, in which private entities were invited to participate if they chose to do so -- and those who did choose to participate did so voluntarily, knowing that the odds of being selected were not high,¹² that the criteria for being included in this experiment were in some measure subjective, and that the goal was the production of experimental information that could be evaluated. Of course, when many more applications were received than anyone had suggested should be prudently introduced at this stage, some evaluation was necessary to attempt to identify those suggestions that might best fit the experimental parameters that had been laid down. But this was never a process in which the absolute or relative merit of the particular application was determinative.

Many applications with likely merit were necessarily not going to be selected, since the goal was a small number (remember, the entire range of responsible suggestions for introducing new TLDs was from one to 10 new ones). And since one objective was diversity -- of business model, of geography, of type of registry -- it was highly likely that some qualified applications would not be selected -- both because prudence required the addition of only a small number of TLDs, and because our proof of concept required data from a diverse set of new TLDs. This was especially true of those applications seeking open, global TLDs; while two were selected, about half of the 44 applications sought such a charter. But it was also true of others; .geo received a very positive evaluation from the staff, but the Board felt that, at this proof of concept stage, there were in fact potential risks to the operation of the DNS that could not be fully evaluated without consultation with the technical support organizations associated with ICANN.

¹² In the application instructions, each applicant was told: "The requirements for sponsoring or operating a new TLD are very stringent. Only a limited number of TLDs will be established in this round of applications, and it is likely that only applications with very high qualifications will be accepted." <http://www.icann.org/tlds/new-tld-application-instructions-15aug00.htm#12>. To make doubly sure there was no misunderstanding, every applicant was required to acknowledge in writing: "The applicant understands and acknowledges that ICANN has the right to reject all applications for new top-level domains that it receives and that there is no assurance that any additional top-level domain will ever be created in the future." <http://www.icann.org/tlds/tld-app-unsponsored-transmittal-15aug00.htm#B6>

Thus, the Board considered every one of the 44 remaining applications at its meeting on November 16, 2000, measuring them against their collective judgment about how well they would serve to carry out the experiment. Although some suggest that the decision process was somehow hidden, in fact all of this consideration was conducted in a public meeting, in full view of the assembled audience and of hundreds of users observing the webcast. In a meeting that lasted more than six hours, the Board methodically reviewed, and either set aside or retained for further evaluation, application after application, until it was left with approximately 10 applications that seemed to have broad consensus support. After further, more focused discussion, that number was pared to the seven that were ultimately selected, and which had almost unanimous Board support: .biz, .info, .pro, .aero, .coop, .museum, and .name.¹³ In the aggregate, the Board concluded that this group provided enough diversity of business models and other relevant considerations so as to form an acceptable test bed or proof of concept.

The various TLDs have very different intended purposes, and that is the strength of the group in the aggregate. Two -- .biz and .info -- were advanced as essentially alternatives to .com -- global, business-oriented registries aimed at capturing millions of registered names around the world. In order to compete with .com -- which has a recognized brand, a large installed base that produces a regular stream of renewals, and a very substantial marketing budget -- these particular applicants assumed they would need a significant investment in both capital equipment and marketing. The Board felt that these applicants seemed most capable of bringing the necessary resources to bear to test whether anyone can effectively compete with .com after the latter's significant head start.

Two other TLDs -- .pro and .name -- were aimed at individuals rather than businesses, but in very different ways. .pro was aimed at licensed professionals, while .name was aimed at any individual. The other three -- .aero (aerospace industry), .coop (for cooperatives), and .museum (for museums) -- were all restricted TLDs, aimed at an industry or a business method or a type of entity, and added to the diversity of this experimental collection of TLDs.

¹³ See <http://www.icann.org/minutes/prelim-report-16nov00.htm#00.89>

ICANN's objectives and, we believe, the objectives of the general Internet community, were to introduce a small number of various kinds of new TLDs into the namespace in a prudent fashion, see what happened, and then, if appropriate, based on those results, move forward with additional new TLDs. It is certainly conceivable that some different subset of the applications it had before it would have met that objective as well as those chosen, but the real question is whether the choices were reasonable, and likely to produce the necessary information on which future introductions could be based. It is also possible, as some of those not selected have complained, that those selected will have a head start (to the extent that matters) over future TLD applicants, but this would be an inevitable consequence of any selection of less than all applicants. Those who were not selected, no matter who they are, were predictably going to be dissatisfied, and those who were selected were predictably going to be glad, but neither was an ICANN goal. ICANN's goal, and its responsibility, was to find a limited collection of diverse new TLDs that could be prudently added to the namespace while minimizing any risk of instability. While time will tell, at this point we believe we faithfully carried out that responsibility.

D.7 The Post-Selection Process. Since November, we have been in the process of drafting and negotiating agreements with the selected applicants. Since these agreements will hopefully be templates for future agreements, we are taking great care to make sure that the structure and terms are replicable in different environments. Since these agreements will contain the promises and commitments under which the applicants will have to live for some time, the applicants are being very careful. The result is slow progress, but progress. We are hopeful that we will be able to complete the first draft agreements within a few weeks. The Board will then be asked to assess whether the agreements reflect the proposals that were selected and, if so, to approve the agreements. Shortly thereafter, this great experiment will begin. We are all looking forward to that time.

Of course, it cannot be stressed enough that no one knows for sure what the effects of this experiment will be. Since there have been no new global TLDs introduced for more than a decade, the Internet is a vastly different space than it was the last time this happened. (There have been a number of country code TLDs introduced over that period, and some of those have recently begun to function in a way quite analogous to a global TLD. These have only achieved relatively small numbers of registrations, so that they do

little to test whether the stability of the DNS by large TLDs competitive with .com.) But there has never been an introduction of as many as seven new global TLDs simultaneously, with the possibility of a land rush that is inherent in that fact. There has never been a highly visible introduction of multiple new TLDs in the context of an Internet that has become a principal global medium for commerce and communication. We do not know whether the introduction of a number of new TLDs -- especially combined with the relatively new phenomenon of the use of ccTLDs in a fashion never intended (after all, .tv stands for Tuvalu, not television, no matter what its marketers say) -- will create consumer confusion, or will impair the functioning of various kinds of software that has been written to assume that .com is the most likely domain for any address.

In short, it is not absolutely clear what effects these introductions will have on the stability of the DNS or how to introduce new TLDs in a way that minimizes harmful side-effects, and that is precisely why we are conducting this experiment. The results will guide our future actions.

IV. IMPORTANT OUTSTANDING ISSUES

A. Country Code Top Level Domain (ccTLD) Relationships

This is certainly one of the most complex parts of the ICANN structural process that remains to be resolved. While there are many moving parts, the key issues are the proper relationships between governments, current ccTLD operators, and ICANN.

To properly understand how we got to where we are, we need to look back to the early days of the DNS, when Jon Postel and others were seeking primarily to expand connectivity throughout the globe. In order to have a truly global network, and for all of the world's population to enjoy the benefits of that network, worldwide connectivity is a crucial first step. After the creation of the original seven generic TLDs in the mid-1980s, Dr. Postel (in his IANA role) delegated what he described as "country code" or "cc" TLDs to persons (often academic researchers) willing to operate those registries for the benefit of the residents of that particular geographic area.

While in general these delegations were made on a national boundary basis, Dr. Postel also made delegations to persons willing to take on this commitment in isolated geographies, such as island groups, even though they might be part of an already existing national cc delegation. Typically, each ccTLD was operated by a designated individual.

Since the goal was to expand connectivity, and since there was in fact very little interest in this subject on the part of most national governments at the time, there was clearly less care and precision about the specifics of those delegations than might seem desirable today.

Over time, the standards and criteria for such delegations grew more rigorous, and were eventually described in a document known as RFC 1591.¹⁴ It became the practice to only create new delegations for those nations or geographic areas included on a list maintained by the Organization for International Standardization (ISO) on behalf of the United Nations Statistics Office, which maintains two-letter codes for nations and various external territories. But the legacy of those early days still remains in some instances, and so there are separate ccTLDs for locations such as the Cocos (Keeling) Islands (.cc) and Christmas Island (.cx), both territories of Australia, as well as various French overseas departments and territories (e.g., Guadeloupe (.gp) and Mayotte (.yt)) and miscellaneous others (e.g., American Samoa (.as) and Pitcairn Island (.pn)).

The 244 ccTLDs are quite diverse. Some, like .de (Germany) or .uk (Great Britain) are large and active registries; some, like .aq (Antarctica), have almost no registrations at all; and some are completely inactive. In addition, the way the ccTLDs are operated varies enormously. Some are highly restricted to residents or citizens of the particular country, while others are completely unrestricted. Some are limited to particular kinds of registrations, while others allow registrations of almost any string of letters. Some are operated as non-profit cooperatives, while others are highly entrepreneurial businesses.

A few delegees have decided to essentially license the marketing of the ccTLD to a commercial enterprise for various forms of compensation, and that has produced out-of-territory marketing campaigns for such ccTLDs as .tv (Tuvalu), .md (Moldova), .nu (Niue), and .cc (Cocos Islands). This practice, of course, is a distortion of the original intended use of the ccTLDs by Jon Postel: opening up the Internet to all parts of the globe, allowing it to accommodate diversity in linguistic, cultural, economic, political, and legal circumstances. Dr. Postel was seeking stewards for the local community's interest in being part of this growing global network we now call the Internet. He was looking for, and

¹⁴ "Requests for Comments" (RFCs") are standards and other related documents developed and published under the auspices of the Internet Engineering Task Force.

generally found, volunteers who were willing to take what he (and they, in the vast majority of cases) viewed as a public trust for the Internet community of which the person or entity receiving the delegation was a part.

But Dr. Postel, while a genius in his field, was no more prescient than anyone else about what the Internet would become, and in any event did not insist upon written contracts and legal agreements as a condition of a delegation. He relied, as did all involved at the time, on the good faith and interest in serving the public of all involved. The Internet we see today -- a global medium for commerce and communication that presents enormous opportunities for profit -- is vastly different than the infant network that he was trying to nurture to adolescence. And it is this evolution that is the principal reason that Dr. Postel and the great majority of the Internet community concluded that something like ICANN needed to be created.

Unfortunately, the task that ICANN inherited with respect to the ccTLDs is complicated greatly by the fact that much water has passed under the bridge since most of the original delegations. In some countries, the national government is intimately involved in oversight of the ccTLD delegated for that country. In other countries, the national government has shown little or no interest in these issues. And in some countries or geographic areas, the operation of the ccTLDs have run strongly counter to governmental preferences and, in several cases, legislation. As we have seen, in some of those situations, where a private entrepreneur has obtained an agreement to market a ccTLD from the original delegee, that entrepreneur has interests that are completely unrelated to the original goals of the delegation, and may have a significant economic interest in maintaining the status quo, without regard to the interests of the global or local Internet communities.

One of the great changes over the last two decades is the involvement and interest of many of the world's governments in the Internet. For many reasons, national governments and even multi-national bodies now see the Internet as an important vehicle or tool for economic development, for communication, for cultural preservation or many other objectives. Since ICANN is a private sector entity, governments or government officials cannot directly participate in the governance or operations of ICANN. But since the interests of national governments are obviously important elements of the global

Internet community, and certainly have to be taken account of in the creation of any meaningful global consensus, ICANN created its Governmental Advisory Committee (GAC) to serve as a device for sharing information and concerns between national governments and recognized geographic areas, on the one hand, and the private sector participants in ICANN on the other. The GAC meets (at least) quarterly, in conjunction with the quarterly ICANN meetings, and has developed the tradition of issuing a public communiqué following each such meeting. ICANN's bylaws provide that it will receive advice from the GAC, and deal with it as it sees fit, but any such advice is always given the serious consideration that it deserves, since it comes from those chosen as representatives of populations that make up much of the Internet community.

The GAC has a particular interest in the relationships between ccTLD delegees and ICANN, since it and its members believe that nations have a sovereign interest in the ccTLD that has been created to represent the interests of the citizens of those nations. On the other hand, while ICANN has recognized and accepted that sovereign interest,¹⁵ ICANN's prime directive -- the stability of the Internet -- cannot be sacrificed or ignored because of it. And finally, there are commercial relationships that have been created over the years, and investments made in reliance on those relationships, that should not be ignored either, both as a matter of common equity and because to do so could adversely affect the very stability that ICANN (and the Internet community it represents) seeks. Thus, finding the correct articulation of the appropriate relationships between (1) the government and the ccTLD administrator; (2) the ccTLD administrator and ICANN; and (3) ICANN and the government is inherently complex.

We have worked hard, but we don't have this part of the ICANN construction project done yet. We have just hired the former ccTLD administrator for Austria, Herbert Vitzthum, to help bring this particular set of conversations to a conclusion as soon as possible.

B. At Large (or general user) Participation in ICANN

The original ICANN proposal to the Department of Commerce to be recognized as the Internet community private sector consensus entity called for by the White Paper

¹⁵ The White Paper summarized this interest as follows: "Of course, national governments now have, and will continue to have, authority to manage or establish policy for their own ccTLDs."

proposed that nine of the 19 members of the ICANN Board would be selected with the participation of the general user community, "if feasible." The reason for the qualification was purely practical; there was a clear consensus that users should have some appropriate way to participate in the ICANN consensus development process, but there was no consensus -- and indeed no good ideas -- about how that could be accomplished without risking the very ability of ICANN to carry out its principal mission. There were loud voices calling for global online elections -- one-person (or one-domain name, or one-email address), one vote. But there were other equally loud voices raising concerns about the possibility that such elections could be easily captured by determined minorities -- business, geographic, ethnic, religious, philosophical or other -- and even if that did not happen, having half the Board of a organization whose principal mission was technical stability of the Internet selected by an electorate that almost certainly would know almost nothing about the subject was unwise. No one had come up with a scheme for accomplishing the objective while eliminating the obvious risks, and so the original proposal promised only to do what was feasible.

The Clinton Administration, in the form of Ira Magaziner, who had been delegated the lead responsibility in this area, took the position that it would not recognize ICANN unless it committed to user participation (comprising at least nomination) in the process of selecting nine Board members -- presumably on the basis that, without that commitment, it would not accept that ICANN truly represented a consensus of the Internet community. Whatever the merits of that point of view, the ICANN organizers acquiesced, and thus began 18 months of searching for a method to accomplish this goal that could achieve consensus support. We have not yet found that needle in the haystack, and this failure has been a constant source of criticism from a portion of the Internet community and from various academics and interest groups who feel strongly about the issue. An appropriate solution for the At Large issue is necessary if ICANN is to truly be an effective consensus development body for the entire Internet community.

Because all other attempts had failed, the ICANN Board last year determined that it would, on an interim basis, conduct a global online selection process for five members of the ICANN Board, which would be followed with a new study effort to find the consensus that has eluded the community so far on this issue. Those selections were held last fall,

and selected one Director from each of the five ICANN regions -- North America, Europe, Asia-Pacific, Latin America, and Africa. The reason for selecting five rather than nine was to obtain support for this interim compromise from those who oppose direct online selections; with five directors selected by online voting, those directors in the aggregate (even in the event of a total capture of the selection process by some determined minority) would still amount to one director less than one-third of the Board. The Board would still be able to act, even where a two-thirds majority was required, even if all five (in this hypothetical) "captured" directors voted as a block. To assure that nine At Large seats continued to be occupied, four of the original interim directors agreed to serve beyond their original two year commitment.

In fact, the online selection process was unsatisfactory in a variety of ways, but it does not appear that there was any "capture" of the five selected directors as a group. They have taken their Board seats and are functioning in every way as ICANN directors today. The At Large Study Committee is in the process of being established, with Carl Bildt, the former prime minister of Sweden as its Chair, and Charles Costello, an elections expert from the Carter Center, and Pindar Wong, a well-known Internet technologist from Hong Kong, as Vice-Chairs. The remainder of the Committee will be named soon, and its goal is to find that consensus solution to this problem that has gone undiscovered to date.

V. CONCLUSION

I trust that this description of the background and current status of the most important ICANN initiatives will be helpful to the Committee, and we stand ready to provide any other information that the Committee might find useful.